Performance of Roadrunner under a Realistic Application Workload

Kevin J. Barker, M. Kei Davis, Adolfy Hoisie, Darren J. Kerbyson, Mike Lang, Scott Pakin, Jose C. Sancho, CCS-1

Performance modeling is an important capability developed and refined by the Performance and Architecture Lab (PAL) over the last few years that has been applied to numerous applications and systems at Los Alamos, within the tri-Lab community, and beyond. Using performance models, systems that are not currently installed and hence not available for measurement can be analyzed and their achievable performance can be quantified. Similarly, the performance of applications running on these systems can be analyzed, predicted with high accuracy, and optimized.

Performance modeling has been used throughout the life-cycle of many systems from first design through to maintenance, as shown in Fig. 1. Models have been used to steer early system designs, for instance in collaboration with IBM in the Defense Advanced Research Projects Agency (DARPA) High Productivity Computer Systems (HPCS) program on their Power 7 system (to be available in 2010). Models were used during the procurement of Roadrunner to compare the performance from different system proposals. Models have also been used to analyze systems during installation and, as evident from the ASCI Q installation, they can be used to identify performance issues and assist in system optimizations.

A PAL performance model analyzes both application and system characteristics. Application characteristics are defined uniquely for each application and include processor flow, data structures used, frequency of use and mapping onto the system, and their potential for resource contention. System characteristics include node configuration (processors per node, shared resources), and interprocessor communication (latency, bandwidth, topology). Many of these are measured (for an existing system) or need to be specified/simulated (for a future system). A separate performance model is developed for each application of interest and thus the approach is application-centric.

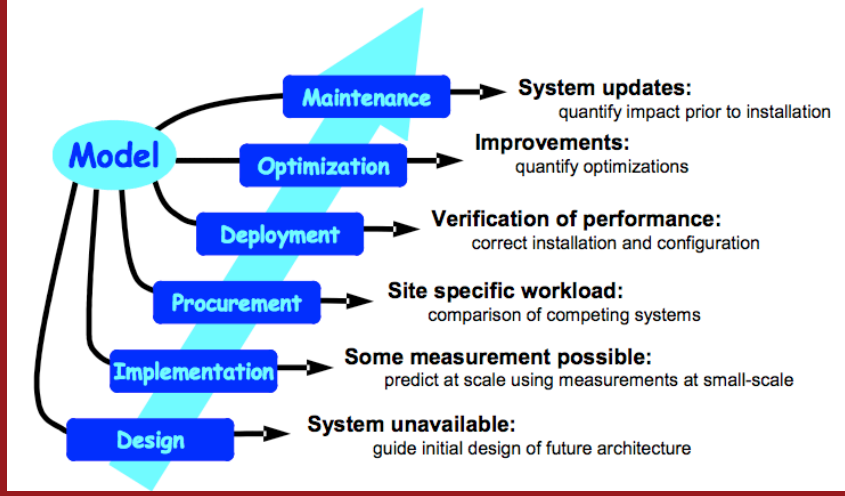

Fig. 1. Performance models have been used throughout the life-cycle of many systems.

The performance of Roadrunner has been analyzed using PAL's highly accurate modeling techniques. PAL developed models of four applications to analyze the performance of Roadrunner prior to its configuration being finalized, and prior to its installation at Los Alamos. This aided in the finalization of the system configuration, and in providing expected performance at large processor-counts for assessment activities. Optimized versions of several Los Alamos applications were used in this process including: Sweep3D (deterministic transport), VPIC (particle-in-cell), SPaSM (molecular dynamics), and Milagro (nondeterministic transport). Sweep3D for the IBM Cell Broadband Engine (Cell BE) was developed by PAL, while the others by researchers in CCS-2.

Many different performance studies were undertaken in analyzing the performance of Roadrunner. For instance, a comparison between the use of Roadrunner in hybrid mode (using both the Cell BEs and Opterons) vs using only the Opterons is shown in Fig. 2 when using only a single node, a single CU (180 nodes), and all CUs. The performance improvement in hybrid mode is over six times for VPIC, over five times for Milagro, approximately four times for Sweep3D on full system and over nine times on a single-node, and over two times for SPaSM. It should be noted in this

performance comparison that the optimization versions of the codes are still being refined and hence the performance improvements may be higher by the time of system installation at Los Alamos.

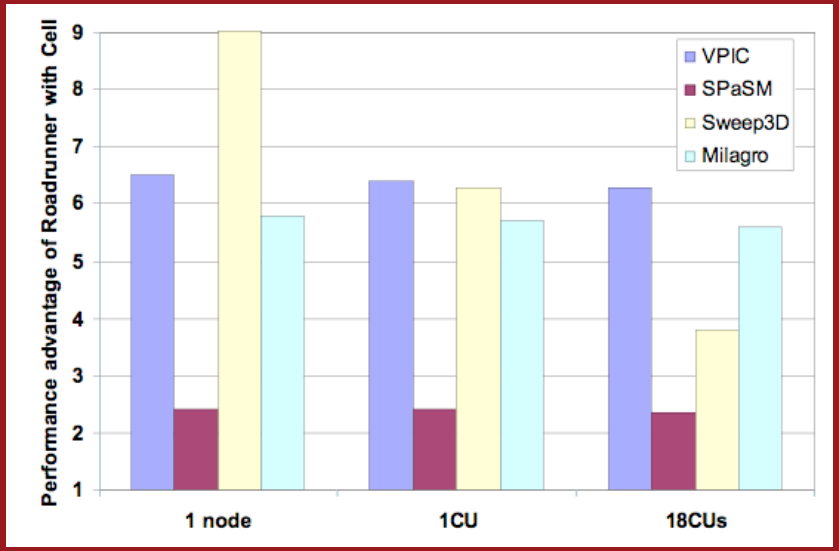

Fig. 2. Performance comparison of using the Roadrunner vs using only AMD Opterons.

Further performance comparisons have shown that the performance of Roadrunner when using the Cell BE will be significantly higher than the performance seen on ASCI Q. A further performance study has shown that the Roadrunner should also achieve higher performance, between 1 and 2.4 times higher, than a system using state-of-the-art AMD or Intel quad-core processors (four sockets per node).

For more information contact Adolfy Hoisie at hoisie@lanl.gov.

References

PAL Roadrunner Report #5, "A Note on the Application Performance of the eDP version of the Cell," Los Alamos National Laboratory report LA-CP-07-1147.

K.J. Barker, et al., "Experiences in Scaling Scientific Applications on Current Generation Quad-core Processors," Los Alamos National Laboratory report LA-UR-07-6855.

Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program